



VisionLabs
MACHINES CAN SEE

LUNA Index Module

Release Notes

Contents

LUNA Index Module v.5.57.0	3
LUNA Index Module v.5.53.0	3
LUNA Index Module v.5.51.6	4
LUNA Index Module v.5.51.4	4
LUNA Index Module v.5.51.0	4
LUNA Index Module v.5.47.4	5
LUNA Index Module v.5.38.1	5
LUNA Index Module v.5.36.5	7
LUNA Index Module v.5.34.0	7

LUNA Index Module v.5.57.0

LIM changes

- API versions for routes of all resources except “/version” and “/healthcheck” have been added to the Indexer service specification.

LIM fixed errors

- Fixed the error where the Python Matcher Proxy service used 100% CPU when using LIM.
- Indexing at startup is now skipped if the time specified in the “planning_period” parameter of the “LIM_MANAGER_INDEXING” section has not passed since the last indexing.

For example, if “planning_period” = 600 seconds (10 minutes) and 5 minutes after the execution of the planning procedure, LIM was restarted, then indexing will not be performed at startup.

- An offline version of the website with online documentation has been added to the LUNA Index Module distribution package, replacing the previous HTML documents.

LUNA Index Module v.5.53.0

Changes

- Starting from the current version of LIM, it is necessary to use LUNA PLATFORM v.5.53.0 and higher. Older versions of LUNA PLATFORM are not supported by the current version of LIM.
- New Redis Sentinel connection settings have been added.

In the “sentinel” section of the “LIM_MANAGER_DB” and “LIM_MATCHER_DB” settings, the “user” and “password” parameters have been added.

To configure “REDIS_URL”, the authorization data must be entered as follows: `redis+sentinel://sentinel_username:sentinel_password@sentinel0:26379,sentinel1:26378/indexed_matcher?username=master_username&password=master_password`.

See the section “Use Redis Sentinel” in the LIM administrator manual for more information about Redis Sentinel.

- A new request “[get list of plugins](#)” has been added to the Indexer, Indexed Matcher and Index Manager services, which enables you to get a list of imported plugins and their status.

Fixed errors

- Fixed the hanging of the Indexed Matcher service when the index refresh is enabled (parameter “lim_matcher_refresh” > “enabled”).

LUNA Index Module v.5.51.6

Changes

- Support for [ECS](#) logging format has been added.

To use the new format, you need to set the value “ecs” in the “format” setting of the “LIM_service_LOGGER” section.

When using the “ecs” value, the following fields will be used:

- “http.response.status_code” - contains the HTTP response status code (e.g., 200, 404, 500, etc.).
 - “http.response.execution_time” - contains information about the time taken to execute the request and receive the response.
 - “http.request.method” - contains the HTTP request method (GET, POST, PUT, etc.).
 - “url.path” - contains the path in the request’s URL.
 - “error.code” - contains the error code if the request results in an error.
- Error descriptions for requests [from the “tasks” section](#) of the Index Manager service have been updated.

Fixed errors

- Fixed the error where Redis settings were reset during LIM settings migration.

LUNA Index Module v.5.51.4

Fixed errors

- Fixed the issue with locking for index search/update causing an “internal server error”.
- Fixed slow matching work on large indexes when index updating is enabled (“LIM_MATCHER_REFRESH” setting) and continuous addition of faces.

LUNA Index Module v.5.51.0

Changes

- The size of the Docker image of lim-configs has been reduced.
- A new setting “db_number” has been added to the sections “LIM_MANAGER_DB” and “LIM_MATCHER_DB”, which enables you to specify the number of the Redis database.

In the setting, you can specify a number from 0 to 15. Each number corresponds to a separate database, which enables you to separate the data.

- SDK updated to version 5.16.0.

Support for the new 57 and 62 neural network model for face descriptor extraction has been added to LUNA Index Module.

LUNA Index Module v.5.47.4

Changes

- The ability to use Redis Sentinel has been added.

New sections “LIM_MANAGER_DB.SENTINEL” and “LIM_MATCHER_DB.SENTINEL” have been added to the settings of the Index Manager and Indexed Matcher services, containing the following settings:

- “master_name” - Name of the Redis master node that will be monitored and managed by Redis Sentinel (by default “index_manager” and “indexed_matcher”).
- “sentinels” - List of addresses of Redis Sentinel listeners (not set by default).

Also, the use of Redis Sentinel is supported in the “REDIS_URL” setting in the “LUNA_INDEXED_LIST_PLUGIN” section, which is responsible for connecting the matching plugin with Redis when calculating matching cost of the request. Example of filling in the “REDIS_URL” setting in the “LUNA_INDEXED_LIST_PLUGIN” section: `redis+sentinel://localhost:26379,localhost:26378/indexed_matcher`.

- The default value for the “CONNECTION_POOL_SIZE” settings of the Index Manager and Indexed Matcher services has been increased from 5 to 100.
- The Python version has been updated to 3.11 in LIM service containers.

Support for older versions of Python has been discontinued.

Fixed errors

- The priority of an error like “Couldn’t find idx for ” has been changed from “ERROR” to “WARNING”.

Such an error can occur when a user has added a new face to the list and then removed it from the list during the period between the launches of index updates in the memory of the Indexed Matcher service (by default, 1 second).

For more information about the index update process, see the “Refreshing index in memory” section of the LIM administrator manual.

LUNA Index Module v.5.38.1

Changes

- Sections of the form `<service_name_HTTP_SETTINGS` have been added to the settings of the Indexer, Index Manager and Indexed Matcher services, containing settings responsible for processing HTTP connections.

The following settings are available:

- `request_timeout` - the duration of time between the instant when a new open TCP connection is passed to the server. Value (in seconds) is integer number, default 60.
- `response_timeout` - the duration of time between the instant the server passes the HTTP request to the app, and the instant a HTTP response is sent to the client. Value (in seconds) is integer number, default 600.
- `request_max_size` - how big a request may be (bytes). Value (in bytes) is integer number, default 1gb.
- `keep_alive_timeout` - http keep alive timeout. Value (in seconds) is integer number, default 15.

See the following link for details: <https://sanic.dev/en/guide/deployment/configuration.html#builtin-values>

- New `socket_address` tag has been added to the index build series monitoring points. The tag contains a value in the format `<host>:<port>`.

This enables you to fix the server address for a particular instance of Indexed Matcher for monitoring.

- New `duration` field has been added to index build series monitoring points with the stages `drop_index` (removing the index from the Indexed Matcher memory) and `load_index` (loading the index into the Indexed Matcher memory). This field indicates the duration of the corresponding stage.
- New resource `"/indexes"` has been added to the Indexed Matcher service, which enables you to get a set of indexes loaded into the service's memory. For each index, its meta information is output ("label", "index_id", "account_id", descriptor version, etc.).
- New `FORMAT` setting has been added to all LIM services, which enables you to specify the logging format - `default` (standard log output format) and `json` (json log output).
- Now if the index building process is interrupted due to lack of memory, the error `26109 Build process killed. It's probably due to running out of memory and the OS was triggering the OOM killer` will be displayed in the logs of the Indexer service.
- Indexed Matcher service now includes a check for the existence of a list upon launching, loading a new index into memory, and updating the index in memory. An index without an existing list will be removed from the service's memory.
- Matching requests have been speeded up.

The speedup was achieved through internal optimizations and changing the response field result format. It used to be a list of the following `{"face_id": ..., "similarity": ...}` objects. The result is now a list of pairs, where the first element of the pair is `face_id` and the second element is `similarity`.

Note. This change requires a mandatory update of the Python Matcher Proxy service.

- Sequence diagrams of the main processes of the LUNA Index Module have been added to the administrator manual.

LUNA Index Module v.5.36.5

Changes

- Performance of matching on a huge index was improved with some internal optimizations.
- In all LIM services, logging to a file is disabled by default (setting `log_to_file` of each service).

LUNA Index Module v.5.34.0

Changes

- The LUNA Index Module installation manual has been updated.
Now by default there are commands to install the module in `/var/lib/luna/` directory instead of in `/var/lib/luna/current/` directory.
The index storage is now created in the `/var/lib/luna/` directory to simplify the upgrade process.
- The LUNA Index Module [upgrade manual](#).
- LUNA Index Module now takes into account changes in the version of face descriptors. It means that:
 - the Indexer service builds an index from descriptors of the version specified in the `DEFAULT_FACE_DESCRIPTOR_VERSION` setting of the Index Manager service;
 - the Index Manager service automatically rebuilds the index if it does not contain information about versions of descriptor;
 - the Indexed Matcher service loads only those indexes that contain descriptors of the version specified in the `DEFAULT_FACE_DESCRIPTOR_VERSION` setting of the Index Manager service.

In this regard, the mandatory field `descriptor_version` has been added to the index metastructure (`meta.json`). The `get indexes` and `get most relevant indexes` requests also return the `descriptor_version` parameter.

Important information for upgrading from previous version

After starting the Index Manager service, it will automatically start rebuilding all indexes that do not contain information about descriptors, i.e. all indexes created in a previous version of LIM. Rebuilding the index can take a long time, depending on the number of faces on the lists. In order to avoid the lengthy process of rebuilding the index, you can add the “descriptor_version” field with the corresponding version of the descriptors to the “meta.json” files of all previously created indexes before starting the Index Manager service.

Reminder added to [“Important information”](#) section of the new LIM upgrade manual.

- Now, by default, the Indexed Matcher service monitors changes in lists with faces. If new changes are made to the list, the Indexed Matcher service updates the corresponding indexes in its memory by gradually adding a small number of descriptors.

The use of this functionality is controlled by the “enabled” setting of the “LIM_MATCHER_REFRESH” section of the Configurator service.

When a cached index is updated, the Indexed Matcher service stops matching on that index, but continues to accept new match requests for that index. By adding a small number of descriptors to the cached index, the matching process is performed with minimal interruption. However, it should be taken into account that if elements are inserted into the list too often (dozens and hundreds of additions), then this can cause a significant degradation in the speed of work, up to an almost complete stop of the matching process.

If this functionality is used, then it is not necessary and not recommended to perform frequent index rebuilds. Accordingly, it is recommended to increase the planning routine period (“planning_period” setting of the “LIM_MANAGER_INDEXING” section of the Configurator service). However, adding new faces to the cached index is slower than rebuilding the index, so it makes no sense to use this function if a very large number of faces have been added to the list. In this case, it is easier to rebuild the index again.

See [“Cached index refreshing”](#) section in the LIM administrator manual.

- Downloading descriptors from the Faces service has been optimized. The “Accept” header has been changed from `application/x-flutbuf` to `application/x-msgpack`.
- “Warming up” (test matching) the first created index before starting to use it has been added. Subsequent indexes are not “warmed up”.

“Warm-up” is performed after the index is loaded into the memory of the Indexed Matcher service and the service is waiting for a matching request.

- The Python version has been updated to 3.10. Support for other versions has been discontinued.